Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	17049	(707/20? 707/2 709/203 709/213 709/214 709/215 709/226 711/147 711/152 711/153 711/114).ccls.	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2005/02/04 09:27
L2	1280	L1 and (fragment$7 or defragment$7)	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2005/02/04 09:27
L3	270	L2 and (((file or object) with (access$3 or use or used or usage or "i/o" or activity) with (frequen$4 or most or greatest or highest)) or lru or lfu or mru or mfu or lfuda or mfuda)	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2005/02/04 09:27
L4	98	L3 and ((disk or storage or system or server or process$3 or distribut$4 or balanc$3 or thread) near4 (idle or off?peak or "off peak" or background or ((low or little) near1 activity)))	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2005/02/04 09:28
L5	59	L4 and @ad<="20011129"	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2005/02/04 09:28
S1	21	((disk or storage or system or processor) near2 (idle or inactive or (low near1 activity))) same (file near3 (maintenance or move or moved or moving or archiv$3 or defragment$7))	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2005/02/03 16:44
S2	1526	(lru or lfu or lfuda or ((recent$4 or frequen$4) near1 (use or used or access$2))) and ((disk or storage or server) with (load or use or usage or activity) with (balanc$3 or distribut$3))	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2005/02/03 16:53
S3	513	S2 and ((idle or inactive or background or off?peak or ((low$2 or less or little) near1 activity)) with (process$3 or distribut$3 or balanc$3 or cycl$3 or maintenance or defragment$8 or moved or moving or movement))	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2005/02/03 17:03
S4	158	S3 and ((file or block or segment) with (frequen$4 or most or least or greatest or amount) with (access$3 or "i/o" or activity or use or used or usage))	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2005/02/03 17:06

S5	182	S3 and ((file or block or segment or object) with (frequen$4 or most or least or greatest or amount) with (access$3 or "i/o" or activity or use or used or usage))	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2005/02/03 17:06
S6	11	S5 and (((file or object) near2 fragment$7) or defragment$7)	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2005/02/03 17:20
S7	107	S5 and (fragment$7 or defragment$7)	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2005/02/03 17:22
S8	70	S7 and @ad<="20011129"	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2005/02/03 17:22
S9	17049	(707/20? 707/2 709/203 709/213 709/214 709/215 709/226 711/147 711/152 711/153 711/114).ccls.	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2005/02/03 17:45
S10	1280	S9 and (fragment$7 or defragment$7)	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2005/02/03 17:46
S11	270	S10 and (((file or object) with (access$3 or use or used or usage or "i/o" or activity) with (frequen$4 or most or greatest or highest)) or lru or lfu or mru or mfu or lfuda or mfuda)	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2005/02/04 09:23

PORTAL
US Patent & Trademark Office

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

+"file system" +disk +background +balanc* +fragment*

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used **file system disk background balanc fragment** Found **68** of **150,138**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 68 Result page: **1** 2 3 4 next

Relevance scale

1 Data partitioning and load balancing in parallel disk systems
Peter Scheuermann, Gerhard Weikum, Peter Zabback
February 1998 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 7 Issue 1
Full text available: pdf(310.27 KB) Additional Information: full citation, abstract, index terms

Parallel disk systems provide opportunities for exploiting I/O parallelism in two possible ways, namely via inter-request and intra-request parallelism. In this paper, we discuss the main issues in performance tuning of such systems, namely striping and load balancing, and show their relationship to response time and throughput. We outline the main components of an intelligent, self-reliant file system that aims to optimize striping by taking into account the requirements of the applications, an ...

Keywords: Data allocation, Disk cooling, File striping, Load balancing, Parallel disk systems, Performance tuning

2 Manageability, availability, and performance in porcupine: a highly scalable, cluster-based mail service
Yasushi Saito, Brian N. Bershad, Henry M. Levy
August 2000 **ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 3
Full text available: pdf(2.52 MB) Additional Information: full citation, abstract, references, index terms

This paper describes the motivation, design and performance of Porcupine, a scalable mail server. The goal of Porcupine is to provide a highly available and scalable electronic mail service using a large cluster of commodity PCs. We designed Porcupine to be easy to manage by emphasizing dynamic load balancing, automatic configuration, and graceful degradation in the presence of failures. Key to the system's manageability, availability, and performance is that sessions, data, and underlying ...

Keywords: cluster, distributed systems, email, group membership protocol, load balancing, replication

3 Serverless network file systems
T. E. Anderson, M. D. Dahlin, J. M. Neefe, D. A. Patterson, D. S. Roselli, R. Y. Wang
December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles**, Volume 29 Issue 5

Full text available: pdf(2.48 MB) Additional Information: full citation, references, citings, index terms

4 Serverless network file systems

Thomas E. Anderson, Michael D. Dahlin, Jeanna M. Neefe, David A. Patterson, Drew S. Roselli, Randolph Y. Wang

February 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 1

Full text available: pdf(2.69 MB) Additional Information: full citation, abstract, references, citings, index terms

We propose a new paradigm for network file system design: serverless network file systems. While traditional network file systems rely on a central server machine, a serverless system utilizes workstations cooperating as peers to provide all file system services. Any machine in the system can store, cache, or control any block of data. Our approach uses this location independence, in combination with fast local area networks, to provide better performance and scalability th ...

Keywords: RAID, log cleaning, log structured, log-based striping, logging, redundant data storage, scalable performance

5 Manageability, availability and performance in Porcupine: a highly scalable, cluster-based mail service

Yasushi Saito, Brian N. Bershad, Henry M. Levy

December 1999 **ACM SIGOPS Operating Systems Review , Proceedings of the seventeenth ACM symposium on Operating systems principles**, Volume 33 Issue 5

Full text available: pdf(1.62 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper describes the motivation, design, and performance of Porcupine, a scalable mail server. The goal of Porcupine is to provide a highly available and scalable electronic mail service using a large cluster of commodity PCs. We designed Porcupine to be easy to manage by emphasizing dynamic load balancing, automatic configuration, and graceful degradation in the presence of failures. Key to the system's manageability, availability, and performance is that sessions, data, and underlying serv ...

6 Petal: distributed virtual disks

Edward K. Lee, Chandramohan A. Thekkath

September 1996 **Proceedings of the seventh international conference on Architectural support for programming languages and operating systems**, Volume 31 , 30 Issue 9 , 5

Full text available: pdf(1.10 MB) Additional Information: full citation, abstract, references, citings, index terms

The ideal storage system is globally accessible, always available, provides unlimited performance and capacity for a large number of clients, and requires no management. This paper describes the design, implementation, and performance of Petal, a system that attempts to approximate this ideal in practice through a novel combination of features. Petal consists of a collection of network-connected servers that cooperatively manage a pool of physical disks. To a Petal client, this collection appear ...

7 Distributed file systems: concepts and examples

Eliezer Levy, Abraham Silberschatz

December 1990 **ACM Computing Surveys (CSUR)**, Volume 22 Issue 4

Full text available: Additional Information: full citation, abstract, references, citings, index

pdf(5.33 MB) terms, review

The purpose of a distributed file system (DFS) is to allow users of physically distributed computers to share data and storage resources by using a common file system. A typical configuration for a DFS is a collection of workstations and mainframes connected by a local area network (LAN). A DFS is implemented as part of the operating system of each of the connected computers. This paper establishes a viewpoint that emphasizes the dispersed structure and decentralization of both data and con ...

8 Frangipani: a scalable distributed file system
Chandramohan A. Thekkath, Timothy Mann, Edward K. Lee
October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles**, Volume 31 Issue 5
Full text available: pdf(2.20 MB) Additional Information: full citation, references, citings, index terms

9 File and storage systems: The Google file system
Sanjay Ghemawat, Howard Gobioff, Shun-Tak Leung
October 2003 **Proceedings of the nineteenth ACM symposium on Operating systems principles**
Full text available: pdf(275.54 KB) Additional Information: full citation, references, index terms

Keywords: clustered storage, data storage, fault tolerance, scalability

10 RAID: high-performance, reliable secondary storage
Peter M. Chen, Edward K. Lee, Garth A. Gibson, Randy H. Katz, David A. Patterson
June 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 2
Full text available: pdf(3.60 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Disk arrays were proposed in the 1980s as a way to use parallelism between multiple disks to improve aggregate I/O performance. Today they appear in the product lines of most major computer manufacturers. This article gives a comprehensive overview of disk arrays and provides a framework in which to organize current and future work. First, the article introduces disk technology and reviews the driving forces that have popularized disk arrays: performance and reliability. It discusses the tw ...

Keywords: RAID, disk array, parallel I/O, redundancy, storage, striping

11 HFS: a performance-oriented flexible file system based on building-block compositions
Orran Krieger, Michael Stumm
August 1997 **ACM Transactions on Computer Systems (TOCS)**, Volume 15 Issue 3
Full text available: pdf(383.67 KB) Additional Information: full citation, abstract, references, citings, index terms, review

The Hurricane File System (HFS) is designed for (potentially large-scale) shared-memory multiprocessors. Its architecture is based on the principle that, in order to maximize performance for applications with diverse requirements, a file system must support a wide variety of file structures, file system policies, and I/O interfaces. Files in HFS are implemented using simple building blocks composed in potentially complex ways. This approach yields great flexibility, allowing an application ...

Keywords: customization, data partitioning, data replication, flexibility, parallel computing, parallel file system

12 Decentralized storage systems: Taming aggressive replication in the Pangaea wide-area file system

Yasushi Saito, Christos Karamanolis, Magnus Karlsson, Mallik Mahalingam

December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue SI

Full text available: pdf(1.93 MB) Additional Information: full citation, abstract, references

Pangaea is a wide-area file system that supports data sharing among a community of widely distributed users. It is built on a symmetrically decentralized infrastructure that consists of commodity computers provided by the end users. Computers act autonomously to serve data to their local users. When possible, they exchange data with nearby peers to improve the system's overall performance, availability, and network economy. This approach is realized by aggressively creating a replica of a file w ...

13 Hints for computer system design

Butler W. Lampson

October 1983 **ACM SIGOPS Operating Systems Review , Proceedings of the ninth ACM symposium on Operating systems principles**, Volume 17 Issue 5

Full text available: pdf(1.73 MB) Additional Information: full citation, abstract, references, citings, index terms

Experience with the design and implementation of a number of computer systems, and study of many other systems, has led to some general hints for system design which are described here. They are illustrated by a number of examples, ranging from hardware such as the Alto and the Dorado to applications programs such as Bravo and Star.

14 The ITC distributed file system: principles and design

M. Satyanarayanan, John H. Howard, David A. Nichols, Robert N. Sidebotham, Alfred Z. Spector, Michael J. West

December 1985 **ACM SIGOPS Operating Systems Review , Proceedings of the tenth ACM symposium on Operating systems principles**, Volume 19 Issue 5

Full text available: pdf(1.01 MB) Additional Information: full citation, references, citings, index terms

15 Computing curricula 2001

September 2001 **Journal on Educational Resources in Computing (JERIC)**

Full text available: pdf(613.63 KB) html(2.78 KB) Additional Information: full citation, references, citings, index terms

16 Process migration

September 2000 **ACM Computing Surveys (CSUR)**, Volume 32 Issue 3

Full text available: pdf(1.24 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Process migration is the act of transferring a process between two machines. It enables dynamic load distribution, fault resilience, eased system administration, and data access locality. Despite these goals and ongoing research efforts, migration has not achieved widespread use. With the increasing deployment of distributed systems in general, and distributed operating systems in particular, process migration is again receiving more attention in both research and product development. As hi ...

Keywords: distributed operating systems, distributed systems, load distribution, process migration

17 The Vesta parallel file system



Peter F. Corbett, Dror G. Feitelson
August 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 3

Full text available: pdf(649.08 KB) Additional Information: full citation, abstract, references, citings, index terms, review

The Vesta parallel file system is designed to provide parallel file access to application programs running on multicomputers with parallel I/O subsystems. Vesta uses a new abstraction of files: a file is not a sequence of bytes, but rather it can be partitioned into multiple disjoint sequences that are accessed in parallel. The partitioning—which can also be changed dynamically—reduces the need for synchronization and coordination during the access. Some control over the layout ...

Keywords: data partitioning, parallel computing, parallel file system

18 The evolution of Coda



M. Satyanarayanan
May 2002 **ACM Transactions on Computer Systems (TOCS)**, Volume 20 Issue 2

Full text available: pdf(441.35 KB) Additional Information: full citation, abstract, references, citings, index terms

Failure-resilient, scalable, and secure read-write access to shared information by mobile and static users over wireless and wired networks is a fundamental computing challenge. In this article, we describe how the Coda file system has evolved to meet this challenge through the development of mechanisms for server replication, disconnected operation, adaptive use of weak connectivity, isolation-only transactions, translucent caching, and opportunistic exploitation of hardware surrogates. For eac ...

Keywords: Adaptation, Linux, UNIX, Windows, caching, conflict resolution, continuous data access, data staging, disaster recovery, disconnected operation, failure, high availability, hoarding, intermittent networks, isolation-only transactions, low-bandwidth networks, mobile computing, optimistic replica control, server replication, translucent cache management, weakly connected operation

19 Performance aspects of disk space management



James O. Dyal, Michael K. Draughn
January 1981 **Proceedings of the 13th conference on Winter simulation - Volume 1**

Full text available: pdf(551.59 KB) Additional Information: full citation, abstract, references, index terms

Dynamic Disk Space Management is a set of strategies for allocating files to disk storage devices so as to reduce seek time, device contention, and storage fragmentation. In this paper, the impact on system throughput and user response time is projected, via simulation model results, for a small general-purpose operating system which manages large capacity disk drives.

20 Client-server computing in mobile environments



Jin Jing, Abdelsalam Sumi Helal, Ahmed Elmagarmid
June 1999 **ACM Computing Surveys (CSUR)**, Volume 31 Issue 2

Additional Information: full citation, abstract, references, citings, index

Full text available: pdf(253.31 KB) terms, review

Recent advances in wireless data networking and portable information appliances have engendered a new paradigm of computing, called mobile computing, in which users carrying portable devices have access to data and information services regardless of their physical location or movement behavior. In the meantime, research addressing information access in mobile environments has proliferated. In this survey, we provide a concrete framework and categorization of the various way ...

Keywords: application adaptation, cache invalidation, caching, client/server, data dissemination, disconnected operation, mobile applications, mobile client/server, mobile compuing, mobile data, mobility awareness, survey, system application

Results 1 - 20 of 68 Result page: **1** 2 3 4 next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.
Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player

IEEE Xplore™ ◄ Back to Previous Page

Results Key:
JNL = Journal or Magazine **CNF** = Conference **STD** = Standard

1 **Intelligent, adaptive file system policy selection**
Madhyastha, T.M.; Reed, D.A.;
Frontiers of Massively Parallel Computing, 1996. Proceedings 'Frontiers '96'., Sixth Symposium on the , 27-31 Oct. 1996
Pages:172 - 179

IEEE CNF

2 **Parallel file system experience on the IBM 9076**
Bidwell, B.; Chi, M.; Colao, R.; Herr, B.; Kavaky, J.; Richards, C.;
Mass Storage Systems, 1994. 'Towards Distributed Storage and Data Management Systems.' First International Symposium. Proceedings., Thirteenth IEEE Symposium on , 12-16 June 1994
Pages:131

IEEE CNF

3 **Prefetching with multiple disks for external mergesort: simulation and analysis**
Pai, V.S.; Varman, P.J.;
Data Engineering, 1992. Proceedings. Eighth International Conference on , 2-3 Feb. 1992
Pages:273 - 282

IEEE CNF

4 **A subsystem design approach to continuous-time input-output performance of decentralized multirate sampled-data systems**
Ito, H.; Ohmori, H.; Sano, A.;
Control, 1994. Control '94. Volume 2., International Conference on , 21-24 Mar 1994
Pages:1023 - 1028 vol.2

IEE CNF

5 **UCFS - a novel User-space, high performance, Customized File System for Web proxy servers**
Jun Wang; Rui Min; Yingwu Zhu; Yiming Hu;
Computers, IEEE Transactions on , Volume: 51 , Issue: 9 , Sept. 2002
Pages:1056 - 1073

IEEE JNL

6 **On the "dimensionality curse" and the "self-similarity blessing"**
Korn, F.; Pagel, B.-U.; Faloutsos, C.;
Knowledge and Data Engineering, IEEE Transactions on , Volume: 13 , Issue: 1 , Jan.-Feb. 2001

Pages:96 - 111

IEEE JNL

7 **Head-positioning control using resonant modes in hard disk drives**
Atsumi, T.; Arisaka, T.; Shimizu, T.; Masuda, H.;
Advanced Motion Control, 2004. AMC '04. The 8th IEEE International Workshop on , 25-28 March 2004
Pages:311 - 316

IEEE CNF

8 **Decentralized control approach to multi-rate sampled-data systems**
Azad, A.M.; Hesketh, T.;
Control, Automation, Robotics and Vision, 2002. ICARCV 2002. 7th International Conference on , Volume: 1 , 2-5 Dec. 2002
Pages:468 - 473 vol.1

IEEE CNF

9 **The SBS: an efficient index structure for spatial database applications**
Al-Badameh, A.; Fotouhi, F.;
Database and Expert Systems Applications, 2000. Proceedings. 11th International Workshop on , 4-8 Sept. 2000
Pages:856 - 860

IEEE CNF

10 **Application-specific file prefetching for multimedia programs**
Mitra, T.; Chuan-Kai Yang; Tzi-Cker Chiueh;
Multimedia and Expo, 2000. ICME 2000. 2000 IEEE International Conference on , Volume: 1 , 30 July-2 Aug. 2000
Pages:459 - 462 vol.1

IEEE CNF

11 **Deflating the dimensionality curse using multiple fractal dimensions**
Pagel, B.-U.; Korn, F.; Faloutsos, C.;
Data Engineering, 2000. Proceedings. 16th International Conference on , 29 Feb.-3 March 2000
Pages:589 - 598

IEEE CNF

12 **Pseudo desired plant approach to the synthesis of almost decoupling control**
Jieh-Shian Young; Yo-Jer Lin; Yue-Wei Huang;
Decision and Control, 2002, Proceedings of the 41st IEEE Conference on , Volume: 4 , 10-13 Dec. 2002
Pages:4539 - 4540 vol.4

IEEE CNF

13 **Using semantic information to guide efficient parallel I/O on clusters**

Schulz, M.;
High Performance Distributed Computing, 2002. HPDC-11 2002. Proceedings. 11th IEEE International Symposium on , 23-26 July 2002
Pages:135 - 142

IEEE CNF

14 **Enhancing data migration performance via parallel data compression**
Jonghyun Lee; Winslett, M.; Xiaosong Ma; Shengke Yu;
Parallel and Distributed Processing Symposium., Proceedings International, IPDPS 2002, Abstracts and CD-ROM , 15-19 April 2002
Pages:444 - 451

IEEE CNF

15 **Compiler-directed I/O optimization**
Kandemir, M.; Choudhary, A.;
Parallel and Distributed Processing Symposium., Proceedings International, IPDPS 2002, Abstracts and CD-ROM , 15-19 April 2002
Pages:19 - 24

IEEE CNF
